

**CLAIMS:**

What is claimed is:

- 1 1. A method for managing data in a memory device having  
2 a plurality of blocks, comprising the steps of:  
3 maintaining a first vector having block entries  
4 sorted in order of number of overall block modifications  
5 for each block of the plurality of blocks;  
6 maintaining a second vector having block entries  
7 sorted in order of number of block modifications since a  
8 previous wear-leveling event; and  
9 using the first vector and the second vector to  
10 determine which of the plurality of blocks should have  
11 its associated data relocated to another block.
- 1 2. The method of Claim 1, wherein the first vector is  
2 sorted in descending order of overall block usage and the  
3 second vector is sorted in ascending order of block usage  
4 since the previous wear level event.
- 1 3. The method of Claim 1, wherein the first vector is  
2 sorted in ascending order of overall block usage and the  
3 second vector is sorted in descending order of block  
4 usage since the previous wear level event.
- 1 4. The method of Claim 1, wherein the block entries for  
2 the first and second vectors each comprise a pointer to a  
3 block descriptor for each of the plurality of blocks.

1 5. The method of Claim 4, wherein each said block  
2 descriptor maintains a modification count for its  
3 respective block.

1 6. The method of Claim 5, wherein the modification  
2 count comprises a count of total overall modifications (n  
3 count) for its respective block and total modifications  
4 since a previous wear-leveling event ( $\Delta n$  count) for the  
5 respective block.

1 7. The method of Claim 6, wherein the n count for a  
2 first given block in the first vector is compared to the  
3 n count for a second given block in the second vector to  
4 determine which of the first and second blocks is more  
5 physically worn, and wherein the  $\Delta n$  count for the first  
6 given block in the first vector is compared to the  $\Delta n$   
7 count for the second given block in the second vector to  
8 determine which of the first and second blocks is more  
9 active, and further comprising the step of swapping  
10 contents of the first given block with the second given  
11 block if either of the first given block and second given  
12 block are both the more physically worn block and the  
13 more active block.

1 8. The method of Claim 1, wherein the using step  
2 comprises the step of copying data contained in a block  
3 having higher usage to a block having lower usage.

1 9. The method of Claim 8, further comprising the step  
2 of copying data contained in a block having lower usage  
3 to a block having higher usage.

1 10. The method of Claim 1 wherein the block  
2 modifications are at least one of erasing blocks and  
3 writing blocks.

1 11. The method of Claim 1, wherein the step of using the  
2 first vector and the second vector comprises the steps  
3 of:

4 determining which of a first given block associated  
5 with an entry in the first vector and a second given  
6 block associated with an entry in the second vector is  
7 more physically worn;

8 determining which of the first given block and the  
9 second given block is more active; and

10 swapping contents of the first given block with  
11 contents of the second given block if either of the first  
12 given block and the second given block are both the more  
13 physically worn block and the more active block.

1 12. The method of Claim 11, wherein the more physically  
2 worn determining step and the more active determining  
3 step are repeated for a plurality of blocks associated  
4 with entries in the first and second vectors, and wherein  
5 the swapping step is deferred until each of the plurality  
6 of blocks associated with each of the first and second  
7 vectors have been processed.

1 13. The method of Claim 1, wherein the using step  
2 comprises the steps of:  
3 (i) determining which of a first given block  
4 associated with an entry in the first vector and a second  
5 given block associated with an entry in the second vector  
6 is more physically worn;  
7 (ii) determining which of the first given block and  
8 the second given block is more active;  
9 (iii) if either one of the first given block and the  
10 second given block are both the more physically worn  
11 block and the more active block, updating a swap table to  
12 indicate that contents of the first given block should be  
13 swapped with contents of the second given block;  
14 (iv) repeating steps (i) - (iii) for each block  
15 entry in at least the first vector;  
16 (v) re-sorting the second vector such that the  
17 blocks associated with the block entries contained  
18 therein are sorted in order of number of block  
19 modifications since a previous wear-leveling event;  
20 (vi) repeating steps (i) - (iv) for the re-sorted  
21 second vector; and  
22 (vii) swapping blocks according to the swap table.

1 14. The method of Claim 13, wherein in step (v) the  
2 second vector is sorted in reverse order from the order  
3 in which it was originally sorted in and, in addition,  
4 the first vector is also re-sorted in reverse order from  
5 the order in which it was originally sorted in.

1 15. A system for managing data in a memory device having  
2 a plurality of blocks, comprising:  
3 a first vector having block entries sorted in order  
4 of number of overall block modifications for each block;  
5 a second vector having block entries sorted in order  
6 of number of block modifications since a previous wear-  
7 leveling event; and  
8 means for using the first vector and the second  
9 vector to determine which of the plurality of blocks  
10 should have its associated data relocated to another  
11 block.

1 16. The system of Claim 15, wherein the first vector is  
2 sorted in descending order of overall block usage and the  
3 second vector is sorted in ascending order of block usage  
4 since the previous wear-leveling event.

1 17. The system of Claim 15, wherein the first vector is  
2 sorted in ascending order of overall block usage and the  
3 second vector is sorted in descending order of block  
4 usage since the previous wear-leveling event.

1 18. The system of Claim 15, wherein the first and second  
2 vectors each comprise a pointer to a block descriptor for  
3 each of the plurality of blocks.

1 19. The system of Claim 18, wherein each said block  
2 descriptor maintains a modification count for each of the  
3 plurality of blocks.

1 20. The system of Claim 19, wherein the modification  
2 count comprises a count of total overall modifications (n  
3 count) for its respective block and total modifications  
4 since a previous wear-leveling event ( $\Delta n$  count) for the  
5 respective block.

1 21. The system of Claim 20, wherein the n count for a  
2 first given block in the first vector is compared to the  
3 n count for a second given block in the second vector to  
4 determine which of the first and second blocks is more  
5 physically worn, and wherein the  $\Delta n$  count for the first  
6 given block in the first vector is compared to the  $\Delta n$   
7 count for the second given block in the second vector to  
8 determine which of the first and second blocks is more  
9 active, and further comprising the step of swapping  
10 contents of the first given block with the second given  
11 block if either of the first given block and second given  
12 block are both the more physically worn block and the  
13 more active block.

1 22. The system of Claim 15, further comprising means for  
2 copying data contained in a block having higher usage to  
3 a block having lower usage.

1 23. The system of Claim 22, further comprising means for  
2 copying data contained in a block having lower usage to a  
3 block having higher usage.

1 24. The system of Claim 15 wherein the block  
2 modifications are at least one of erasing blocks and  
3 writing blocks.

1 25. The system of Claim 15, wherein the means for using  
2 comprises:

3 first means for determining which of a first given  
4 block associated with an entry in the first vector and a  
5 second given block associated with an entry in the second  
6 vector is more physically worn;

7 second means for determining which of the first  
8 given block and the second given block is more active;  
9 and

10 means for swapping contents of the first given block  
11 with contents of the second given block if either of the  
12 first given block and the second given block are both the  
13 more physically worn block and the more active block.

1 26. A data storage subsystem comprising a memory  
2 controller, system memory and a plurality of flash  
3 devices, each flash device organized as a plurality of  
4 blocks, wherein the memory controller operates to perform  
5 the steps of:

6 maintaining a first vector having block entries  
7 sorted in order of number of overall block modifications  
8 for each block of the plurality of blocks;

9 maintaining a second vector having block entries  
10 sorted in order of number of block modifications since a  
11 previous wear-leveling event; and

12        using the first vector and the second vector to  
13        determine which of the plurality of blocks should have  
14        its data located to another block.

1        27. A method for managing data in a memory device having  
2        a plurality of blocks, comprising the steps of:  
3                maintaining a first vector having block entries for  
4        each block of the plurality of blocks;  
5                maintaining a second vector having block entries for  
6        each block of the plurality of blocks; and  
7                using the first vector and the second vector to  
8        determine which of the plurality of blocks should have  
9        its associated data relocated to another block.

1        28. The method of Claim 27, wherein a given block entry  
2        in the first vector is associated with a first block of a  
3        block pair and a given block entry in the second vector  
4        is associated with a second block of the block pair, and  
5        wherein the step of using the first vector and the second  
6        vector comprises the step of traversing through at least  
7        a portion of the first vector and the second vector to  
8        compare an overall usage count of a given first block and  
9        a given second block of an associated block pair to  
10       determine which block of the associated block pair is  
11       more overall heavily used.

1        29. The method of Claim 28 wherein the step of  
2        traversing further comprises comparing a recent usage  
3        count of the given first block and the given second block



4 of the associated block pair to determine which block of  
5 the associated block pair is recently more heavily used.

1 30. The method of Claim 29, wherein if the more overall  
2 heavily used block of the associated block pair is the  
3 same as the recently more heavily used block of the  
4 associated block pair, swapping data contents between  
5 each block of the associated block pair.

1

1 31. The method of Claim 29, wherein if the more overall  
2 heavily used block of the associated block pair is not  
3 the same as the recently more heavily used block of the  
4 associated block pair, modifying the given entry for one  
5 of the first vector and second vector to point to a  
6 different given entry, and repeating the overall heavily  
7 used determination and the recently more heavily used  
8 determination.

1 32. The method of Claim 29, wherein if (i) the more  
2 overall heavily used block of the associated block pair  
3 is the same as the recently more heavily used block of  
4 the associated block pair, and (ii) a difference between  
5 overall usage count of the given first block and the  
6 given second block exceeds a predetermined value in  
7 absolute value, and (iii) a difference between recent  
8 usage count of the given first block and the given second  
9 block exceeds a second predetermined value in absolute  
10 value, swapping data contents between each block of the  
11 associated block pair.

1 33. The method of Claim 29, wherein if (i) the more  
2 overall heavily used block of the associated block pair  
3 is not the same as the recently more heavily used block  
4 of the associated block pair, or (ii) the difference  
5 between overall usage count of the given first block and  
6 the given second block does not exceed a predetermined  
7 value in absolute value, or (iii) the difference between  
8 recent usage count of the given first block and the given  
9 second block does not exceed a second predetermined value  
10 in absolute value, modifying the given entry for one of  
11 the first vector and second vector to point to a  
12 different given entry, and repeating the overall heavily  
13 used determination, the recently more heavily used  
14 determination, the determination of whether the  
15 difference between overall usage count of the given first  
16 block and the given second block exceeds a predetermined  
17 value in absolute value, and the determination of whether  
18 the difference between recent usage count of the given  
19 first block and the given second block exceeds a second  
20 predetermined value in absolute value.

1 34. The method of Claim 27, wherein a current entry in  
2 the first vector points to a first block of a block pair  
3 and a current entry in the second vector points to a  
4 second block of the block pair, and wherein the step of  
5 using the first vector and the second vector comprises  
6 the step of comparing the overall block usage for a block  
7 associated with the current entry in the first vector  
8 with the overall block usage for a block associated with

9 the current entry in the second vector to determine which  
10 block is more overall heavily used.

1 35. The method of Claim 34, further comprising the step  
2 of comparing a recent usage count of the block associated  
3 with the current entry in the first vector with a recent  
4 usage count of the block associated with the current  
5 entry in the second vector to determine which block is  
6 recently more heavily used

1 36. The method of Claim 35, wherein if the more overall  
2 heavily used block is the same as the recently more  
3 heavily used block, swapping data contents of the block  
4 associated with the current entry in the first vector  
5 with data contents of the block associated with the  
6 current entry in the second vector.

1 37. The method of Claim 35, wherein if the more overall  
2 heavily used block is not the same as the recently more  
3 heavily used block, modifying the current entry for one  
4 of the first vector and second vector and repeating the  
5 overall heavily used determination and the recently  
6 heavily used determination.

1 38. The method of Claim 35, wherein if (i) the more  
2 overall heavily used block is the same as the recently  
3 more heavily used block, and (ii) a difference between  
4 overall usage count of the block associated with the  
5 current entry in the first vector and the block  
6 associated with the current entry in the second vector

7 exceeds a predetermined value in absolute value, and  
8 (iii) a difference between recent usage count of the  
9 block associated with the current entry in the first  
10 vector and the block associated with the current entry in  
11 the second vector exceeds a second predetermined value in  
12 absolute value, swapping data contents of the block  
13 associated with the current entry in the first vector  
14 with data contents of the block associated with the  
15 current entry in the second vector.

1 39. The method of Claim 35, wherein if (i) the more  
2 overall heavily used block is not the same as the  
3 recently more heavily used block, or (ii) a difference  
4 between overall usage count of the block associated with  
5 the current entry in the first vector and the block  
6 associated with the current entry in the second vector  
7 does not exceed a predetermined value in absolute value,  
8 or (iii) a difference between recent usage count of the  
9 block associated with the current entry in the first  
10 vector and the block associated with the current entry in  
11 the second vector does not exceed a second predetermined  
12 value in absolute value, modifying the current entry for  
13 one of the first vector and second vector to point to a  
14 different current entry, and repeating the overall  
15 heavily used determination, the recently more heavily  
16 used determination, the determination of whether the  
17 difference between overall usage count of the block  
18 associated with the current entry in the first vector and  
19 the block associated with the current entry in the second  
20 vector exceeds a predetermined value in absolute value,

Docket No. STK-2003-038-SFT

21 and the determination of whether the difference between  
22 recent usage count of the block associated with the  
23 current entry in the first vector and the block  
24 associated with the current entry in the second vector  
25 exceeds a second predetermined value in absolute value.